Predicting Public Health Impacts of Electricity Usage

Anonymous Author(s)
Affiliation
Address
email

Abstract

Electricity consumption impacts public health due to pollutant emissions from fossil fuel power plants. While stricter regulations have reduced emissions, fossil fuels remain a dominant energy source, necessitating advanced methods to quantify and mitigate these societal health effects. Here, we present a domain-specific AI model, `HealthPredictor`, an end-to-end pipeline that links electricity usage to public health outcomes. Our system integrates three key components: a fuel mix predictor that forecasts energy source contributions, an air quality converter that models pollutant emissions and dispersion, and a health impact assessor that translates environmental changes into monetary health costs. We demonstrate that our health-driven optimization approach achieves significantly lower prediction errors compared to fuel mix-driven methods across multiple U.S. regions. Through a case study on electric vehicle charging schedules, we show the public health benefit of our approach in providing actionable insights about electricity usage for users. This work thus demonstrates how AI model can be explicitly designed to optimize for public health and societal well-being.

1 Introduction

Electricity generation significantly impacts public health through the emission of air pollutants, yet quantifying and predicting societal health impacts remains challenging due to the complex relationships between electricity usage, emissions, pollutant dispersion, and health outcomes [15, 6]. The urgency of understanding these relationships has intensified with the rapid growth of large energy loads. For instance, the rise of artificial intelligence and large language models has led to unprecedented energy demand from data centers [2]. This trend, combined with the increasing electrification of transportation and industrial processes, makes electricity usage a critical sector for mitigating public health impacts, a critical topic of social well-being.

Electricity consumption directly impacts public health through air pollution from fossil fuel power plants, which remain the largest industrial polluters [37, 42]. This relationship between electricity use and health impacts offers unique opportunities for intervention because electricity demand is dynamically *controllable*, unlike other natural pollution sources or weather patterns. This controllability enables proactive demand-side management by tapping into energy load flexibilities, e.g., scheduling data center workloads or coordinating electric vehicle charging schedules in residential sectors.

Moreover, electricity generation is a significant and growing pollution source, with health impacts extending far cross-state [37]. Despite strict U.S. regulations reducing power sector emissions, fossil fuel plants continue to be "a leading source of air, water, and land pollution that affects communities nationwide," as reported by the EPA [37]. Analysis using the EPA's COBRA modeling tool [34] indicates that health costs from electricity are on track to rise, rivaling those of on-road emissions by 2028 as shown in Figure 1.

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

Coal-fired power plants are among the fossil fuel facilities with the most adverse health effects, with their $PM_{2.5}$ emissions estimated to have caused approximately 460,000 excess deaths between 1999 and 2020 [11]. Despite their significant health impact, they remain a key component of the U.S. electricity mix. Importantly, the U.S. EIA projects that even by 2050, fossil fuels will still account for a significant share of electricity generation, with coal power generation remaining around 180 billion kWh under the alternative scenario where power plants are allowed to operate subject to rules existing before early 2024 [33]. The continued reliance on fossil fuels means that even though the U.S. is among the leading countries in clean energy development, the associated health risks cannot be overlooked. Therefore, it is crucial to predict and mitigate these health impacts through targeted interventions alongside supply-side grid decarbonization.

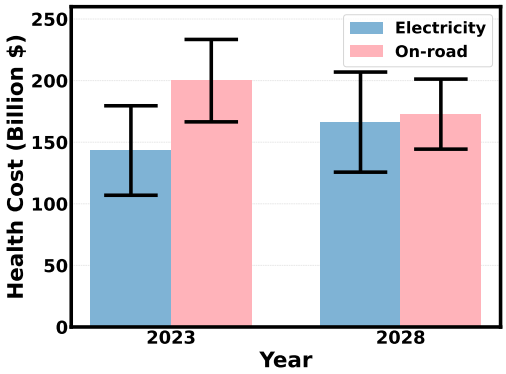

Figure 1: Total public health costs of electricity generation and on-road emissions in the contiguous U.S. in 2023 and 2028 [38]. The error bars represent high and low estimates provided by COBRA using two different exposure-response models.

Prior research evolves from epidemiological pollution-health correlations to advanced ML for modeling complex environmental systems. Early studies quantified the health effects of air pollutants, particularly particulate matter, laying the foundation for air dispersion models like COBRA [34] and InMAP [27] by leveraging models like Gaussian dispersion equations and chemical transport simulations to estimate pollutant spread and associated health risks. Recent advances in machine learning have transformed this research landscape. [25] demonstrated the effectiveness of LSTM networks in air quality assessment and pollution forecasting, achieving higher accuracy than traditional statistical methods. Researchers have developed foundation models that integrate diverse data sources to forecast comprehensive atmospheric composition [1].

While these advances are significant, most existing research addresses isolated aspects of the problem, either assessing health impacts from air pollutants or modeling energy-to-emissions conversion [6, 1]. Given that fossil fuel generation will remain a substantial part of the power grid for the foreseeable future, there is a critical gap to design a domain-specific model capable of connecting demand-side usage directly to health outcomes. Such predictions would provide valuable signals to users, enabling them to take informed actions to mitigate air pollution to protect public health.

We present a domain-specific AI model, `HealthPredictor`, an end-to-end pipeline that quantifies the public health impacts of electricity consumption patterns. Our system integrates three key components: a fuel mix predictor that forecasts the proportional contribution of different energy sources to electricity generation, an air quality converter that models pollutant emissions and their atmospheric dispersion, and a health impact assessor that translates environmental changes into monetary health costs. By combining these components with a health-driven optimization framework, `HealthPredictor` enables prediction of health impacts from electricity usage patterns. We demonstrate the effectiveness of our approach through an electric vehicle (EV) charging case study, where users can determine optimal charging schedules to minimize adverse health outcomes. In addition, we release the datasets we have collected and processed to help advance the field of research by addressing the limitations of fragmented and dispersed data from various sources, observed in previous works [1]. Our approach bridges the critical gap between electricity consumption and health outcomes, providing actionable insights for both individuals and system operators.

2 Related Works

Energy system modeling typically focuses on technological and economic characteristics, often incorporating health damage in an aggregated and simplified manner [26, 17]. However, these approaches rarely provide granular insights into the direct health impacts of electricity generation. Some methodologies focus on optimizing energy systems to reduce emissions, but their health impact assessments tend to remain indirect or high-level, missing the opportunity for detailed, localized health assessments [19, 15].

Epidemiological studies have made significant contributions to understanding the relationship between health and air pollution [37]. For example, [9] has conducted a comprehensive regional impact assessment of air quality improvement, while [3] developed log-linear models for quantifying asthma hospitalizations based on particulate matter levels, demonstrating the direct correlations between air quality and health outcomes in urban environments. Although these studies provide a valuable foundation for linking environmental pollutants with health statistics, they lack an integrated framework that connects energy systems directly to health outcomes. Air pollution dispersion modeling also plays a critical role in supporting these epidemiological studies. [22] systematically reviewed computational fluid dynamics approaches for urban air pollution modeling, and [27] has developed InMap, a specialized model for analyzing air pollution interventions, accounting for complex atmospheric chemical interactions. While these models offer insights into the dispersion of pollutants, they typically do not link air pollutants to human activities, and thus do not provide actionable insights for system operators or individuals to make improvements.

Recent advances have integrated machine learning with environmental health research. [25] explored the use of LSTM network for air quality assessment and pollution forecasting, demonstrating the potential of data-driven approaches. Additionally, the emerging field of *health-informed computing*, exemplified by works like [10], seeks to quantify the broader societal impacts of technological systems, providing a methodological foundation for the future research on the health consequences of electricity generation, mainly w.r.t. the advancement of AI and the development of large data centers.

Most existing research focuses on individual stages of health impact analysis—either assessing health impacts from air pollutants or modeling energy-to-emissions conversion. Few studies have developed a comprehensive pipeline that directly quantifies health impacts from electricity usage across residential or industrial sectors. While some researchers have proposed holistic frameworks, they often focus on specific fuel sources, such as coal, or are limited to particular regions [6], and they focus on highlighting the importance of the problem with a high level of analysis rather than modelling. To the best of our knowledge, advanced modeling approaches that provide high-accuracy, end-to-end predictions of health impacts from electricity consumption remain scarce. These gaps underscore the need for more comprehensive approaches that encompass various fuel sources and regions, offering actionable insights for system operators and individuals.

3 Background and Problem Formulation

In this section, we review the background and introduce formulations related to health impact assessments from the use of power generation fuel mix.

3.1 Air Pollutants for Health Impact

Air quality, a critical determinant of human health, is shaped by the presence of specific gases and particulate matter in the atmosphere. Six pollutants are recognized as primary contributors to air quality degradation: carbon monoxide (CO), nitrogen oxide (NO), nitrogen dioxide (NO_2), sulfur dioxide (SO_2), ozone (O_3), and particulate matter in three size categories, PM_1, $PM_{2.5}$, and PM_{10} [42]. These pollutants originate from various sources, including fossil fuel combustion and industrial processes [1]. The long-distance transport of these pollutants amplifies their public health impact, particularly for vulnerable populations such as the elderly and individuals with preexisting conditions. Adverse health outcomes, including premature mortality, asthma exacerbation, and cognitive decline, also result in substantial societal costs through increased hospitalizations and medication use [7, 40].

3.2 Air Dispersion Formulation

Establishing meaningful relationships between emission sources and their health impacts is non-trivial. Typically, the first step is to determine the spatial and temporal distribution of pollutants in the area. This process usually involves mathematical models with varying spatial resolutions that solve the governing dispersion-advection equations. By integrating emission data with meteorological inputs, dispersion models can estimate pollutant concentrations at specific receptor points [22].

Assume there are K types of air pollutants and M receptor regions of interest. Let $\mathcal{P}_s = (\mathcal{P}_{s,1}, \ldots, \mathcal{P}_{s,K})$ denote the quantities of K types of air pollutants at the emission source. For receptor i, the corresponding quantities are represented by $\mathcal{P}_r^i = (\mathcal{P}_r^{i,1}, \ldots, \mathcal{P}_r^{i,K})$. A general disper-

sion model can be formulated as:

$$\mathcal{P}_r^1, \ldots, \mathcal{P}_r^M = D_{\boldsymbol{w}}(\mathcal{P}_s), \tag{1}$$

which gives the amount of K types of air pollutants at receptor region $i = 1, \ldots, M$, i.e., $\mathcal{P}_r^i = (\mathcal{P}_r^{i,1}, \ldots, \mathcal{P}_r^{i,K})$. The parameter $\boldsymbol{w}$ consider factors such as geographical conditions, characteristics of emission source, and meteorological data [10, 28].

Despite rapid advancements in mathematical models, the uncertainty still exists due to the complex interplay between emission sources and meteorological conditions [22, 4]. While emission models require detailed anthropogenic data, meteorological predictions depend on both measurements and simulations to capture atmospheric turbulence. These complexities make deterministic air-dispersion models less reliable for forecasting purposes.

3.3 Measuring Health Impacts

The relationship between changes in adverse health effects and changes in air pollution exposure can be quantified using epidemiological studies [37]. For example, the rate of asthma hospitalizations can be modeled as a log-linear function of particulate matter levels [3]. Specifically, for a receptor region i, the change in the number of adverse health effects ΔY^i can be expressed as:

$$\Delta Y^i = Y_0^i \times \text{POP}^i \times \left(1 - e^{-\beta \Delta \mathcal{P}_r^i}\right), \tag{2}$$

where Y_0^i is the baseline incidence rate for the health outcome at receptor i, POP^i is the population exposed at the receptor region, β is the concentration-response coefficient derived from epidemiological studies, and $\Delta \mathcal{P}_r^i = (\Delta \mathcal{P}_r^{i,1}, \ldots, \Delta \mathcal{P}_r^{i,K})$ is the change in pollutant concentrations at receptor i.

3.4 Converting Health Impacts into Monetary Valuation

Health Impact Assessment (HIA) often requires converting health outcomes into monetary values enable cost-benefit analysis and facilitate policy decision-making [29, 37]. We denote these values by $v^i = (v^{i,1}, \ldots, v^{i,H})$, where H represents the number of different types of health impacts (e.g., premature mortality, asthma attacks) at receptor i. Commonly used methodologies for this conversion include estimating the economic value of a statistical life (VSL), as proposed by the Organization for Economic Cooperation and Development (OECD) [21], and quality-adjusted life years (QALYs) [9, 16].

Our work aims to establish the connection between energy consumption from electricity usage at a source s over time steps $t = 1, \ldots, T$, denoted as $E_t^s = (E_1^{s,1}, \ldots, E_T^{s,F})$, where F represents the number of different fuel mix sources (such as oil and gas), and the resulting economic health outcomes v_t^i at receptor i at time step t.

4 Methods

Our methodology integrates diverse datasets and modeling approaches into a cohesive, end-to-end pipeline, which links electricity consumption to public health outcomes based on the power generation fuel mix pattern, named the *HealthPredictor*. As shown in Figure 2, the framework consists of three core modules: the Fuel Mix Predictor, the Air Quality Converter, and the Health Impacter.

4.1 Modelling Framework

4.1.1 Fuel Mix Predictor

Fuel mix predictor is the starting point of our pipeline, as they directly inform the potential health impacts of electricity generation. Fossil fuels (e.g. coal, gas, oil) are particularly significant due to their high pollutant emissions. Electricity demand also affects the types of fuel used for power generation. For example, during peak electricity demand periods, less-efficient and higher-emission fuel sources like coal or oil are often used [13]. Variations in the fuel mix, driven by factors such as demand, renewable availability, and regulations, cause fluctuations in pollutant emissions, which in

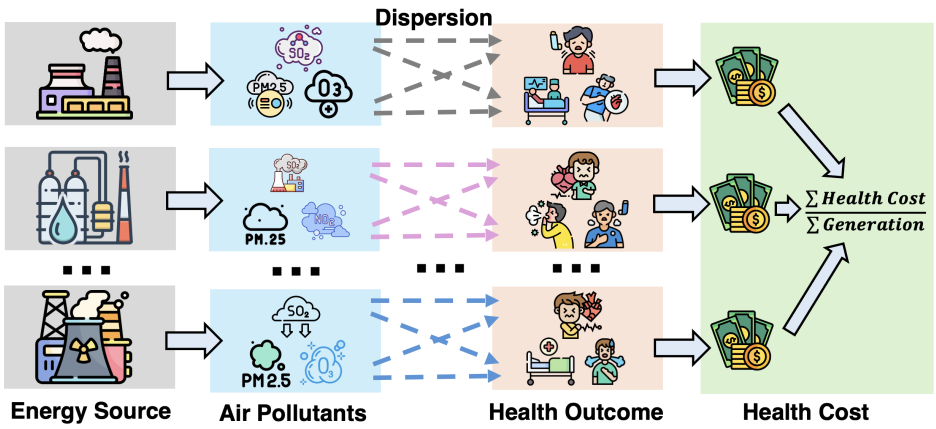

Figure 2: Overview of the health-informed computing pipeline. The pipeline begins with energy contribution, E'_t, from various sources (e.g. gas, coal). It then models pollutant dispersion (e.g. $S0_2$, NO_x) to receptors. Finally, it quantifies the resulting health impacts by monetary cost metrics ($/MWh).

turn impact human health outcomes like respiratory illnesses and premature deaths [1]. Therefore, accurately predicting the fuel mix is crucial to estimate these emissions and assess their potential health impacts.

The goal of the fuel mix predictor is to estimate the future fuel mix (e.g., coal, oil, gas) utilized for electricity generation across the next time horizon based on historical data on the grid's fuel mix. The fuel mix data is inherently time-variant [18], and accurately predicting it requires a deep understanding of the dynamic interplay between factors such as electricity demand, market conditions, and regulations. Several machine learning approaches have been proposed to model these chaotic and nonlinear time-series relationships, including the LSTM networks to capture temporal dependencies in energy forecasting [12]. Hybrid approaches, such as combining wavelet basis functions (WBF), sparse autoencoders (SAE), and LSTM, also aim to improve prediction accuracy by integrating multiple advanced [23]. For our predictor, we opt for the Transformer-based architecture in favor of its superior ability to capture long-range dependencies, given the lengthy yearly dataset we have collected.

4.1.2 Air Quality Converter

The Air Quality Converter is tasked with transforming predicted fuel mix data into quantifiable air pollutant emissions in our pipeline.

Pollutants Estimation Each fuel type has distinct emission factors that determine the amount of pollutants emitted per unit of energy generated. These factors can vary depending on combustion technology, fuel quality, and operating conditions. Additionally, regional differences, such as local fuel types, regulatory standards, and emission control technologies, can further influence these emission factors [37]. By obtaining the fuel mix predictions from the fuel mix predictor, we can estimate pollutant emissions based on these factors using tools provided by InMap [27] or COBRA [34].

Dispersion Modelling Modelling the dispersion of air pollutants is a critical step in understanding the relationship between emissions and their resulting concentrations in the atmosphere. This process provides insight into how pollutants spread, dilute, and interact with environmental conditions, ultimately determining their impact on air quality and public health. Two primary types of pollutants are usually considered in dispersion modeling. One is the primary pollutants, such as directly emitted particulate matter (e.g., PM2.5). It usually exhibit a linear relationship with source emissions. These pollutants can be effectively modeled using tools like AERMOD, a steady-state Gaussian plume dispersion model recommended by the EPA [37], which calculates pollutant concentrations by accounting for environmental variables such as wind speed, atmospheric stability, and source

characteristics. The other is the secondary pollutants, such as O_3 formed from precursors NO and NO_2, involve more complex, non-linear relationships with emissions. Their formation results from chemical reactions in the atmosphere, influenced by factors like sunlight and temperature. To model these interactions, chemical transport models CTMs) that simulate atmospheric chemistry and transport processes are commonly used [27].

As we introduced in Section 3.2, the general relationship between pollutant concentrations at receptor sites and emissions from sources can be expressed by Eq. (1). In this formulation, $D_{\boldsymbol{w}}(\mathcal{P}_s)$ encapsulates the complex dispersion process, where $\boldsymbol{w}$ accounts for factors such as geographical conditions, emission source characteristics, and meteorological data. In our framework, the dispersion function $D_{\boldsymbol{w}}(\mathcal{P}_s)$ is modeled using a custom neural network layer designed to approximate the transformation $f(E_s)$, which represents the transformation applied to the emissions E_s. This neural network layer takes as input the pollutant quantities at the emission source ($\mathcal{P}_s$) along with relevant environmental features, and predicts pollutant quantities at the receptor regions ($\mathcal{P}_r^i$ for $i = 1, \ldots, M$). The neural network parameters, corresponding to $\boldsymbol{w}$, are trained to minimize the discrepancy between predicted and observed concentrations. By leveraging the neural network-based dispersion modeling, we aim to provide a more accurate and flexible framework for predicting pollutant concentrations.

4.1.3 Health Impacter

The health impact component aims to quantify the changes in adverse health effects resulting from variations in air pollution exposure, following the results obtained from the air quality converter. We measure the health impact in dollars per megawatt-hour ($/MWh). This measurement reflects the economic cost of health impacts associated with electricity generation [21]. The most widely utilized functional form in criteria air pollutant concentration-response modeling is the log-linear model, as introduced in Eq. (2). This model is well-suited to capture non-linear relationships between pollutant concentrations and health risks, particularly at lower exposure levels where the relative risk changes more sharply. In addition to the log-linear model, linear models are also applied in specific cases, such as when evaluating the health impacts of certain pollutants (e.g., SO_2 or $PM_{2.5}$) or within specific demographic groups where simpler proportional relationships may better describe the data [37].

The health impact modeling process aligns closely with established epidemiological frameworks and tools such as COBRA [37], which quantify the proportional increase in health risks due to incremental changes in pollutant concentrations. These models incorporate concentration-response functions derived from epidemiological studies, enabling a robust estimation of health risks, such as premature mortality and respiratory illnesses [37]. By leveraging these functions, our pipeline calculates the economic valuation of health impacts per unit of electricity generation.

4.2 End-to-End Training

Loss Function Design We design a loss function for a health-informed learning pipeline with incorporating health impact measures directly into the optimization process.

Let y_t be the true fuel mix at time t and $\hat{y}_t$ be the predicted fuel mix at time t by the fuel mix predictor. In AirQuality converter, $f(\cdot)$ be the function that converts fuel mix predictions to pollutant emissions like SO_2 after dispersion. Let $g(\cdot)$ be the function that estimates health impacts measured by $/mWh based on pollutant levels, and spatial features, denoted as I. In the formulation, we also introduce a hyperparameter β to balance forecasting accuracy and health impact optimization. The loss function is formulated as

$$\mathcal{L} = \beta \left\| y_t - \hat{y}_t) \right\|^2 + (1 - \beta) \left\| y_{\text{impact},t} - g(\hat{y}_t, I) \right\|^2, \tag{3}$$

where y_{impact} denotes the true value of the health impact, while $g(\hat{y}_t, I)$ represents the predicted health impact, based on the predicted fuel mix and a series of models that convert the corresponding pollutants into health impacts, and β is a hyperparameter to balance the prediction accuracy of the fuel mix and the health impact. It is worth noting that $y_{\text{impact},t}$ here is general, which can refer to local health impacts, global health impacts, or a combination of both, depending on the context and the extent of the dispersion of pollutants in the atmosphere. This flexibility allows our pipeline to account for both direct localized effects and broader regional or global consequences of air pollution.

Table 1: Comparison of methods for health impact predictions. The **Health-driven Opt** based on the *transformer architecture in the fuel mix predictor component* achieved the lowest MSE loss on both internal and external health impact predictions compared to other methods, across various prediction time windows and regions.

Region	Method	Prediction Window (hrs)	Fuel-Mix Prediction Loss	**Health Impact Loss (Internal)**	**Health Impact Loss (External)**
CISO	Fuel-mix-driven Opt (LSTM)	$T = 24$	0.1005	1.1417	1.1882
		$T = 72$	0.1181	1.1645	2.2554
	Fuel-mix-driven Opt	$T = 24$	**0.0160**	1.4451	1.7562
		$T = 72$	0.0195	1.5354	2.0057
	Health-driven Opt (LSTM)	$T = 24$	0.6012	0.0207	0.0233
		$T = 72$	0.7054	0.0203	0.0217
	Health-driven Opt	$T = 24$	0.0718	0.0151	0.0194
		$T = 72$	0.0700	**0.0144**	**0.0182**
PJM	Fuel-mix-driven Opt (LSTM)	$T = 24$	0.1883	1.2244	0.9023
		$T = 72$	0.2101	3.1549	1.0721
	Fuel-mix-driven Opt	$T = 24$	**0.0280**	1.1446	0.4388
		$T = 72$	0.0342	3.0022	1.2097
	Health-driven Opt (LSTM)	$T = 24$	0.4962	0.0341	0.0225
		$T = 72$	0.5302	0.0180	0.0102
	Health-driven Opt	$T = 24$	0.1561	0.0313	0.0136
		$T = 72$	0.1695	**0.0152**	**0.0069**
ERCO	Fuel-mix-driven Opt (LSTM)	$T = 24$	0.1160	2.5192	1.0247
		$T = 72$	0.0586	2.2059	1.0912
	Fuel-mix-driven Opt	$T = 24$	**0.0108**	2.4784	1.0571
		$T = 72$	0.0116	2.3926	0.9938
	Health-driven Opt (LSTM)	$T = 24$	0.4122	0.0232	0.0133
		$T = 72$	0.8132	0.0270	0.0118
	Health-driven Opt	$T = 24$	0.0537	**0.0099**	**0.0043**
		$T = 72$	0.0701	0.0199	0.0083

In training, our pipeline learns to predict health impacts from fuel mix data through multiple stages, using a customized loss function specifically tailored to incorporate accurate health impact prediction, as shown in Eq. (3). The input to our pipeline is a sequence of fuel mix data, which represents the distribution of fuel types over time (e.g., hourly fuel mix data for a year). The output is the predicted health impact, expressed as a monetary value per unit of energy ($/mWh).

5 Experiments

In this section, we implement our pipeline and develop methods, including health impact-driven approaches and optimizations, to demonstrate the effectiveness and flexibility of our pipeline. This implementation serves as the foundation for the case study carried out in Section 6, which signal electric vehicle (EV) users to reduce adverse health outcomes during charging.

Datasets Our analysis uses fuel mix data from U.S. Energy Information Administration (EIA) [32] and health impact data ($/MWh) based on estimates from the AVoided Emissions and geneRation Tool (AVERT) from the latest available year [37].

While our dataset covers all 67 U.S. BA, we our experiments on three representative regions—California (CISO), Texas (ERCO), and the Mid-Atlantic (PJM). These regions are carefully selected to reflect diverse characteristics in grid operations, emission profiles, and public health impact patterns. The dataset includes six input features, such as fuel mix percentages and time period, and two output features: internal (within-BA) and external (outside-BA) health impacts. More details and additional empirical results are provided in Appendix A.1.

Model Construction For the fuel-mix predictor, we develop a Transformer-based architecture tailored to fuel mix time-series data, capitalizing on its ability to capture intricate relationships and long-term dependencies . To model the complex, non-linear conversion of emissions to health impacts, we utilize a 3-layer Multi-Layer Perceptron (MLP) . The detailed model architectures and hyperparameters are provided in Appendix A.2.

Implementation Details We consider two predicted outputs: *Health Impact (Internal)* and *Health Impact (External)*, which account for the dispersion of air pollutants beyond their region of origin [37]. The “Internal" captures the total health cost within a BA’s jurisdiction, while the “External" reflects

the cost in all counties outside that domain. The loss function in Eq. (3) for the case study can then be rewritten as

$$\mathcal{L} = \beta \left\| y_t - \hat{y_t}) \right\|^2 + \frac{1-\beta}{2} \left(\left\| y_{i,t} - g(\hat{y}_t, I) \right\|^2 + \left\| y_{e,t} - g(\hat{y}_t, I) \right\|^2 \right) \tag{4}$$

where $y_{i,t}$ represents the within-region health impact, and $y_{e,t}$ represents the external (outside-region) health impact at time t. In our experiments, we observed no clear justification to prioritize internal versus external health impacts for regions. Therefore, to avoid notation clutter, we set the hyperparameter values for within-region and outside-region health impacts to $\frac{1-\beta}{2}$.

In our experiments, predictions are made across time windows (T) of 24 and 72 hours. In addition to our Transformer-based models, we also implement LSTM-based variants of both the Fuel-mix-driven Opt and Health-driven Opt methods as **baselines** for comparison. Further details on data splitting and optimizer hyperparameters are moved to Appendix A.2.

Main Results For our method, we set $\beta = 0.2$, assigning a weight of 0.8 to the health impact component, evenly split between internal and external health impacts at 0.4 each. This setting, referred to as *Health-driven Opt*, prioritizes health outcomes over fuel-mix prediction accuracy. In *Fuel-mix-driven Opt*, we set $\beta = 1$ to optimize only the fuel-mix predictors as one baseline. In Table 1, we present the prediction results for these methods, showing that Health-driven Opt consistently achieves lower loss in health impacts compared to other baselines across different regions and prediction time steps. Among them, those based on the Transformer architecture achieve lower losses in fuel-mix prediction compared to the ones using LSTM, while Health-driven Opt methods, regardless of the architecture used in its fuel-mix predictor, achieves lower loss in both internal and external health impacts compared to the Fuel-mix-driven Opt methods.

6 A Case Study of Health-Aware EV Charging

The increasing adoption of electric vehicles (EVs) highlights concerns about their potential health-related impacts, particularly through emissions associated with electricity generation. Scheduling EV charging strategically can play a critical role in reducing harmful emissions, thereby mitigating public health risks and also supporting power system stability [5].

Our method provides a straightforward, data-driven signal for EV users by predicting health impacts over the next few hours based on electricity usage patterns. These predictions evaluate health impacts caused by electricity usage, expressed in units of $\$/MWh$, guide users to identify optimal charging times, helping to minimize exposure to pollutants such as NO_2 and $PM_{2.5}$. By delivering quantifiable and actionable insights, our approach empowers users to make informed decisions, effectively reducing the health risks associated with electricity usage. In our case study, different charging schedule strategies along with their corresponding numerical results are presented and analyzed.

Setups For an EV denoted as j, let I_j and D_j represent the initial and target state of charge (SoC), respectively. The charging occurs within a time frame starting at s_j and ending at e_j. To optimize this process, we discretize the interval $[s_j, e_j]$ into time slots $\tau = 1, ..., T$ and implement a binary charging scheme $\mathcal{Y}_j$. Each element $y_{j,t}$ in $\mathcal{Y}_j$ is either 1, indicating charging at time t, or 0, indicating no charge (e.g., $\mathcal{Y}_j = [1, 0, ...1]$). The electricity charged at time t for EV j is denoted by $\zeta_{j,t}$. Considering the health impact H_t at time t, the goal to minimize the total charging health impact by determining the optimal charging schedule for EV j can be expressed as follows,

$$\min_{\mathcal{Y}_j} \sum_t \zeta_{j,t} y_{j,t} \cdot H_t, \quad s.t. \quad \sum_t \zeta_{j,t} y_{j,t} + I_j = D_j, \tag{5}$$

where $Y_j = [y_{j,1}, \cdots, y_{j,T}]$.

EV-Charging Datasets We use the publicly available ACN-Data [14], which provides real-time charging details (e.g., arrival/departure times, energy delivered), to estimate power demand and charging rates for EVs in residential areas.

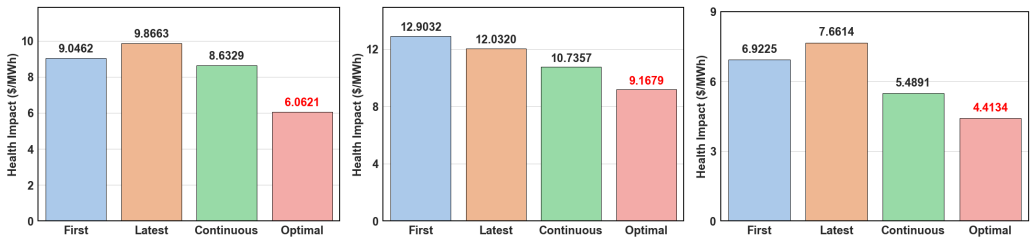

(a) CISO (b) PJM (c) ERCO

Figure 4: Simulation results of using different EV charging strategies based on health impact predictions in CISO, PJM and ERCO regions. With the provided prediction signals from the `HealthPredictor`, EV users can choose the *optimal* hours to charge their vehicles, achieving the greatest adverse health outcomes reduction compared to other charging strategies.

To approximate the available residential charging time window (s_j to e_j), we leverage data from the National Household Travel Survey (NHTS) [30]. We assume the distributions of the initial charging time (s_j) and end time (e_j) align with the NHTS distributions of home arrival and departure times, respectively. For the health impact predictions H_t, we use the empirical results from Section 5 on different regions.

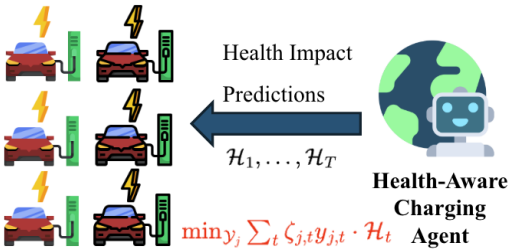

Figure 3: The `HealthPredictor`-embedded health-aware charging agent provides EVs with health outcome estimates for each charging time window, $H_1, \cdots, H_T$. This information allows EVs to determine the optimal charging schedule by minimizing the total health impact, given by $\sum_t \zeta_{j,t} y_{j,t} \cdot H_t$, for each EV j.

Simulation Results We evaluated several charging strategies: *First Hours*, which charges during the earliest available hours after arriving; *Latest Hours*, which charges during the latest available hours before departure; and *Continuous Charging*, which involves charging continuously from an optimal starting time t_1 to satisfy the demand D_j while minimizing the overall health impact.

In the simulation, we use predictions of both internal and external health impacts of Health-driven Opt method from Section 5 to calculate H_t by a weighting factor of 0.5 to balance them. Figure 4 compares the total health impacts generated throughout the entire charging process. By optimizing the charging schedule using Eq. (5) which selects optimal charging hours based on health impact predictions $\mathcal{H}$, significant reductions in total health impacts can be achieved. Specifically, across the CISO, PJM, and ERCO regions, our approach reduces total health impacts by $\sim$24–42% compared to the *First Hours* and *Latest Hours* strategies, and by $\sim$15–20% compared to *Continuous Charging*.

7 Conclusion

This work introduces a novel approach to bridging the gap between electricity consumption decisions and their public health implications. Our `HealthPredictor` demonstrates that incorporating health impact considerations into electricity usage predictions can lead to substantial reductions in adverse health outcomes. The effectiveness of our approach is validated across different U.S. regions and through a practical case study on EV charging optimization, showing potential health impact reductions of 17-42% compared to other charging strategies. By providing quantifiable health impact predictions, our system enables more informed decision-making for both individuals and system operators. Future work could extend this framework by exploring decision-focused learning that optimize the minimization of adverse health outcomes during training.

Limitations. We acknowledge several limitations in our study. For example, our predictions only consider relatively short time windows and do not extend to long-term scenarios. Additionally, while we use the EPA's air dispersion model as the ground truth, there may still be high level of uncertainty in air dispersion due to the complex interplay between emission sources and meteorological conditions [22, 4].

References

[1] Cristian Bodnar, Wessel P. Bruinsma, Ana Lucic, Megan Stanley, Anna Vaughan, Johannes Brandstetter, Patrick Garvan, Maik Riechert, Jonathan A. Weyn, Haiyu Dong, Jayesh K. Gupta, Kit Thambiratnam, Alexander T. Archibald, Chun-Chieh Wu, Elizabeth Heider, Max Welling, Richard E. Turner, and Paris Perdikaris. A foundation model for the earth system. *arXiv preprint*, arXiv:2405.13063, 2024.

[2] Justine Calma. How much electricity does ai consume? *The Verge*, January 2023. Accessed: 2025-01-27.

[3] Kefei Chen, Gary Glonek, Alana Hansen, Susan Williams, Jonathan Tuke, Amy Salter, and Peng Bi. The effects of air pollution on asthma hospital admissions in adelaide, south australia, 2003–2013: time-series and case–crossover analyses. *Clinical & Experimental Allergy*, 46(11):1416–1430, 2016.

[4] Chardin Hoyos Cordova, Manuel Niño Lopez Portocarrero, Rodrigo Salas, Romina Torres, Paulo Canas Rodrigues, and Javier Linkolk López-Gonzales. Air quality assessment and pollution forecasting using artificial neural networks in metropolitan lima-peru. *Scientific Reports*, 11(1):24232, 2021.

[5] Constantin Filote, Raluca-Andreea Felseghi, Maria Simona Raboaca, and Ioan Aşchilean. Environmental impact assessment of green energy systems for power supply of electric vehicle charging station. *International Journal of Energy Research*, 44(13):10471–10494, 2020.

[6] Robert B Finkelman, Amy Wolfe, and Michael S Hendryx. The future environmental and health impacts of coal. *Energy Geoscience*, 2(2):99–112, 2021.

[7] Giulia Grande, Jing Wu, Petter L. S. Ljungman, Massimo Stafoggia, Tom Bellander, and Debora Rizzuto. Long-term exposure to $PM_{2.5}$ and cognitive decline: A longitudinal population-based study. *Journal of Alzheimer's Disease*, 80(2):591–599, 2021.

[8] Gianluca Guidi, Francesca Dominici, Nat Steinsultz, Gabriel Dance, Lucas Henneman, Henry Richardson, Edgar Castro, Falco J Bargagli-Stoffi, and Scott Delaney. The environmental burden of the united states' bitcoin mining boom. 2024.

[9] James Hall, Jian Zhong, Sue Jowett, Andrea Mazzeo, G Neil Thomas, John R Bryson, Steve Dewar, Nadia Inglis, Mark Wolstencroft, Catherine Muller, et al. Regional impact assessment of air quality improvement. *Environmental Pollution*, 356:123871, 2024.

[10] Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, and Shaolei Ren. The unpaid toll: Quantifying the public health impact of ai. *arXiv preprint arXiv:2412.06288*, 2024.

[11] Lucas Henneman, Christine Choirat, Irene Dedoussi, Francesca Dominici, Jessica Roberts, and Corwin Zigler. Mortality risk from united states coal electricity generation. *Science*, 382(6673):941–946, nov 2023.

[12] Seçkin Karasu and Aytaç Altan. Crude oil time series prediction model based on lstm network with chaotic henry gas solubility optimization. *Energy*, 242:122964, 2022.

[13] Christian Klemm and Peter Vennemann. Modeling and optimization of multi-energy systems in mixed-use districts: A review of existing methods and approaches. *Renewable and Sustainable Energy Reviews*, 135:110206, 2021.

[14] Zachary Lee, Tongxin Li, and Steven Low. Acn-data: Analysis and applications of an open ev charging dataset. pages 139–149, 06 2019.

[15] Lixin Li, Yan Lv, Bo Sun, Miao Wang, Bin Chen, Zeke Li, and Haiwei Fan. Integrated energy system evaluation and optimization based on integrated evaluation model and time-series optimization. *Journal of Computational Methods in Sciences and Engineering*, 24(1):19–36, 2024.

[16] James Lomas, Laetitia Schmitt, Sally Jones, Maureen McGeorge, Elizabeth Bates, Mike Holland, Duncan Cooper, Richard Crowther, Mike Ashmore, David Rojas-Rueda, et al. A pharmacoeconomic approach to assessing the costs and benefits of air quality interventions that improve health: a case study. *BMJ open*, 6(6):e010686, 2016.

[17] Roland Cunha Montenegro, Panagiotis Fragkos, Audrey Helen Dobbins, Dorothea Schmid, Steve Pye, and Ulrich Fahl. Beyond the energy system: modeling frameworks depicting distributional impacts for interdisciplinary policy analysis. *Energy Technology*, 9(1):2000668, 2021.

[18] Amir Mosavi, Mohsen Salimi, Sina Faizollahzadeh Ardabili, Timon Rabczuk, Shahaboddin Shamshirband, and Annamaria R Varkonyi-Koczy. State of the art of machine learning models in energy systems, a systematic review. *Energies*, 12(7):1301, 2019.

[19] Esi W Nkyekyer and Andrew L Dannenberg. Use and effectiveness of health impact assessment in the energy and natural resources sector in the united states, 2007–2016. *Impact Assessment and Project Appraisal*, 37(1):17–32, 2019.

[20] North American Electric Reliability Corporation. North american electric reliability corporation, 2025. Accessed: 2025-01-23.

[21] Organisation for Economic Co-operation and Development (OECD). *The Economic Consequences of Outdoor Air Pollution*. OECD Publishing, Paris, 2016.

[22] Mariya Pantusheva, Radostin Mitkov, Petar O Hristov, and Dessislava Petrova-Antonova. Air pollution dispersion modelling in urban environment using cfd: a systematic review. *Atmosphere*, 13(10):1640, 2022.

[23] Weibiao Qiao, Wei Liu, and Enbin Liu. A combination model based on wavelet transform for predicting the difference between monthly natural gas production and consumption of us. *Energy*, 235:121216, 2021.

[24] Jean Schmitt, Marianne Hatzopoulou, Amir FN Abdul-Manan, Heather L MacLean, and I Daniel Posen. Health benefits of us light-duty vehicle electrification: Roles of fleet dynamics, clean electricity, and policy timing. *Proceedings of the National Academy of Sciences*, 121(43):e2320858121, 2024.

[25] Dewen Seng, Qiyan Zhang, Xuefeng Zhang, Guangsen Chen, and Xiyuan Chen. Spatiotemporal prediction of air quality based on lstm neural network. *Alexandria Engineering Journal*, 60(2), 2021.

[26] Avinash Shankar Rammohan Subramanian, Truls Gundersen, and Thomas Alan Adams. Modeling and simulation of energy systems: A review. *Processes*, 6(12):238, 2018.

[27] Christopher W Tessum, Jason D Hill, and Julian D Marshall. Inmap: A model for air pollution interventions. *PloS one*, 12(4):e0176131, 2017.

[28] Christopher W. Tessum, Jason D. Hill, and Julian D. Marshall. Inmap: A model for air pollution interventions. *PloS ONE*, 12(4):e0176131, 2017.

[29] Adrián Giovani Trejo-González, Horacio Riojas-Rodriguez, José Luis Texcalac-Sangrador, Carlos Manuel Guerrero-López, Karla Cervantes-Martinez, Magali Hurtado-Diaz, Luz Angélica de la Sierra-de la Vega, and Pamela Estrellita Zuñiga-Bello. Quantifying health impacts and economic costs of pm 2.5 exposure in mexican cities of the national urban system. *International journal of public health*, 64:561–572, 2019.

[30] U.S. Department of Transportation. Federal Highway Administration. National Household Travel Survey, 2017.

[31] U.S. Energy Information Administration. U.s. energy atlas. Accessed: 2025-01-23.

[32] U.S. Energy Information Administration. Hourly generation fuel mix data, 2023.

[33] U.S. Energy Information Administration. Annual energy outlook 2025 data browser. Web database, April 2025. Reference case projections through 2050.

[34] U.S. Environmental Protection Agency. Co-Benefits Risk Assessment (COBRA) Health Impacts Screening and Mapping Tool. Accessed: 2025-01-23.

[35] U.S. Environmental Protection Agency. Estimating the public health benefits of energy efficiency and renewable energy with epa's benefits per kwh values. Technical report, U.S. Environmental Protection Agency, 2019. Accessed: 2025-04-30.

[36] U.S. Environmental Protection Agency. Emissions & generation resource integrated database (egrid). *EPA eGRID Database*, 2024. Accessed: 2024-10-31.

[37] U.S. Environmental Protection Agency (EPA). *COBRA User Manual v5.1*. U.S. Environmental Protection Agency (EPA), 2024. `https://www.epa.gov/system/files/documents/2024-06/cobra-user-manual-v5.1.pdf`.

[38] U.S. EPA. Co-benefits risk assessment health impacts screening and mapping tool (COBRA). `https://cobra.epa.gov/`.

[39] U.S. EPA. What is cross-state air pollution? `https://www.epa.gov/Cross-State-Air-Pollution/what-cross-state-air-pollution`.

[40] U.S. EPA. Research on health effects from air pollution, 2025. Accessed: 2025-01-22.

[41] U.S. Office of Management and Budget. Circular no. a-4: Regulatory analysis. `https://obamawhitehouse.archives.gov/omb/circulars_a004_a-4`, 2003. Accessed: 2024-01-30.

[42] World Health Organization. *WHO Global Air Quality Guidelines: Particulate Matter (PM2.5 and PM10), Ozone, Nitrogen Dioxide, Sulfur Dioxide and Carbon Monoxide*. World Health Organization, Geneva, 2021.

A Appendix

A.1 Datasets Collection and Preparation Details

We here document details on constructing the comprehensive datasets that link the fuel mix usage of power generation with health outcomes for training the `HealthPredictor`.

Beyond the end-to-end training design, it is important to highlight that the datasets required to train our pipeline are not only complex but also fragmented and labor-intensive to construct. These data come from heterogeneous sources with inconsistent schemas and geographic granularity. For example, aligning fuel categories across agencies (e.g., 8 types in EIA vs. 40 in EPA eGRID) required systematic mapping, and spatial integration between balancing authorities (BAs) and county-level health data involved optimized point-in-polygon indexing. We have created datasets that link hourly fuel mix compositions with corresponding health costs, covering all 67 BAs in the U.S. for the latest available year, totaling 586,920 data points. We release these datasets with our code to support future research and practical applications. More details are reported in the Experiments and Appendix.

Our primary dataset consists of hourly generation fuel mix data and their corresponding internal and external health costs per megawatt-hour (MWh) for the selected geographical regions in the most recent year. To construct this comprehensive dataset, we employed a systematic approach encompassing data acquisition, processing, and analysis through the following procedures:

Step I: Acquisition of Hourly Generation Mix Data We collect the hourly generation fuel mix data of latest years from the U.S. Energy Information Administration (EIA) [32]. The U.S. EIA provides electricity generation data organized by balancing authorities (BAs, a functional role defined by the North American Electric Reliability Corporation [20]) rather than state boundaries. This kind of organization is preferred as BAs align more closely with the operational structure of the power grid, providing a more accurate representation of how electricity is generated and managed across regions. In Figure 5, we report the distribution of different fuel types for the regions we studied: CISO, PJM, and ERCO, averaged hourly throughout the year. For ERCO, the petroleum consumption is processed as 0% in our analysis due to the EIA including it within the broader "Other" category

without specific data. According to [32], petroleum usage in power generation is generally minimal in ERCO, so excluding it as a separate category does not affect the overall fuel mix analysis for our methods. In processing the fuel mix data, missing data points are addressed using a two-step imputation approach. The primary method involves interpolation based on adjacent hourly data. When such data are unavailable, missing values are substituted by averaging corresponding time points from the nearest available days, taking advantage of daily cyclical patterns in the fuel mix.

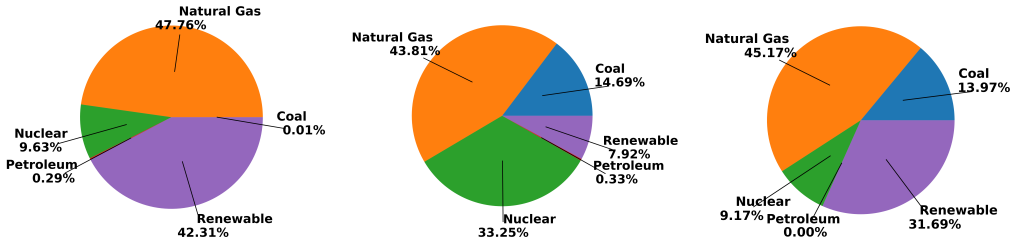

(a) CISO (b) PJM (c) ERCO

Figure 5: Distribution of the energy generation mix by different fuel types in CISO, PJM, and ERCO.

Step II: Derivation of Emission Data Emission data were derived from the Environmental Protection Agency's Emissions & Generation Resource Integrated Database (eGRID) [36], which provides raw plant-level electricity generation and emission data. Our analysis focuses on four criteria air pollutants: $PM_{2.5}$, SO_2, NO_X, and VOC. We obtained pollutant datasets from the most recent available eGRID records: plant-level $PM_{2.5}$ emissions from 2021 and plant-level SO_2, NO_X, and VOC emissions from 2022. Since the eGRID database associates each generation facility with its corresponding balancing authority (BA), we needed to map the raw plant-level data to the specific BAs relevant to our study. For the selected BA, we then assumed a unit electricity consumption (1 MWh) for each hour throughout 2023. Using hourly generation fuel mix data, we allocated this unit hourly demand across different fuel sources according to their generation shares. For each fuel type, we further distributed its allocated generation among all plants within the BA based on their relative generation capacities.

One challenge in this process arose from the inconsistent categorization of fuel types between the EIA and eGRID datasets. To address this, we developed a systematic mapping approach. First, we utilized eGRID's internal hierarchical classification system to map its numerous detailed fuel types to a smaller set of fuel type categories defined within eGRID. Then, we mapped these simplified eGRID categories to EIA's classification system through both direct correspondence (e.g., HYDRO to WAT) and careful examination of category definitions for less straightforward cases (e.g., DFO to OIL). This meticulous mapping process is essential to ensure accurate integration of emissions data with generation profiles. By combining these carefully harmonized plant-level allocations with plant-specific emission factors, we quantified the specific emissions profile per MWh for each hour.

Step III: Health Cost Assessment To assess the public health implications of our emissions profile, we employ the CO-Benefits Risk Assessment (COBRA) Health Impacts Screening and Mapping Tool (Desktop v5.1, as of October 2024) developed by the U.S. EPA [34]. COBRA utilizes a reduced-complexity air quality dispersion model incorporating a source-receptor matrix for expedited assessment. Despite its wide validation and adoption in the literature for large-scale air quality and health impact analyses [24, 8], applying COBRA to derive health costs requires significant effort. It involves labor-intensive steps to compile and prepare the input data, including mapping emissions profiles to specific regions and ensuring the appropriate application of emission factors, all while maintaining the integrity of the tool's assumptions. In derivation, we set 2023 as the baseline scenario year to correspond to our study period. In accordance with EPA recommendations based on the U.S. Office of Management and Budget Circular No. A-4 guidance [41], we implement a discount rate of 2% in the COBRA model.

Considering the air pollutant transport mechanisms, we account for both internal and external health impacts in our analysis of an emission source. The spatial delineation of each BA's service territories is obtained from the U.S. Energy Atlas [31], which provides raw data in GeoJSON format. To efficiently process these complex geographical data, we employed spatial indexing techniques to optimize

Table 2: Comparison of methods for health impact predictions in State Tennessee with $T = 24$.

Method	Fuel-Mix Prediction Loss	**Health Impact Loss (Internal)**	**Health Impact Loss (External)**
Fuel-mix-driven Opt (LSTM)	0.1683	0.2059	0.3096
Fuel-mix-driven Opt	**0.0743**	0.1093	0.6791
Health-driven Opt (LSTM)	0.7288	0.0150	0.2099
Health-driven Opt	0.1524	**0.0135**	**0.1466**

the computational performance of point-in-polygon operations, enabling the precise identification of counties within each BA's operational domain. Following the spatial categorization of counties as either internal or external to the BA, we aggregated the county-level health costs accordingly. Specifically, we compute the hourly internal and external health costs throughout the year based on the unit hourly electricity consumption (1 MWh), where internal costs represent the sum of health impacts in counties within the BA's jurisdiction, and external costs comprise impacts in all other counties. This yields a comprehensive dataset comprising hourly fuel mix compositions and their corresponding internal and external health costs per MWh for the entire year, which is used to train our proposed pipeline.

Brief Summary of Dataset Preparation Challenges Challenges in dataset preparation mainly include reconciling semantic inconsistencies across data sources, resolving spatial mismatches, and managing infrastructure for local health impact computation. For example, the U.S. Energy Information Administration (EIA) defines 8 fuel mix categories, while the EPA's eGRID lists over 40, requiring us to systematically consolidate and map these types using internal hierarchies and cross-referencing definitions. Spatial integration was equally nontrivial—health impacts from COBRA are county-based, whereas eGRID data is organized by BAs. To bridge this gap, we employed U.S. Energy Atlas GeoJSON files with optimized spatial indexing for accurate point-in-polygon assignments. Additionally, estimating county-level health costs using COBRA's desktop tool demanded significant manual effort, with each run taking 5–20 minutes and data entry requiring 1–2 minutes per input. This entire process spanned several days and underscores the substantial effort involved in building a reliable, multi-source dataset.

A.2 Additional Empirical Details and Results

The three regions—CISO (California), PJM (Mid-Atlantic), and ERCO (Texas) selected in our main text have shown the effectiveness of our methods in various energy generation patterns and regulatory environments. These regions are carefully chosen to represent distinct characteristics in power grid operations, emissions profiles, and public health impact patterns. Texas, for example, ranks among the top three for $PM_{2.5}$ emissions, which have severe health effects [39]. Although California does not have the highest emissions, its dense population results in significant adverse health costs [35]. Specifically, CISO has one of the lowest benefits-per-kWh, reflecting high adverse health outcomes, as reported by the EPA [35].

Model Construction For the fuel-mix predictor, we develop a Transformer-based architecture tailored to fuel mix time-series data, capitalizing on its ability to capture intricate relationships between various factors influencing the fuel mix. The transformer also excels at capturing long-term dependencies, which are critical in understanding the temporal dynamics of fuel usage and transitions over extended periods. The architecture consists of an embedding layer followed by a Transformer block with a single encoder and decoder layer, utilizing four multi-head attention mechanisms with a dropout regularization rate of 0.1.

The conversion of pollutant emissions to air pollutant concentrations and their subsequent dispersion in the atmosphere is a highly intricate process. It involves complex chemical transformations, atmospheric reactions, and meteorological processes. To address this complexity, we utilize a 3-layer Multi-Layer Perceptron (MLP) model, which takes the fuel mix predictions as input and predicts the potential health impact. The model is specifically chosen for its ability to approximate complex, nonlinear relationships inherent in pollutant dispersion and their effects.

In experiments, the LSTM based fuel mix predictor is composed of an embedding layer that projects inputs to a 64-dimensional space, followed by a single-layer LSTM with 64 hidden units and a dropout rate of 0.1. The number of training epochs is set to 150 for Transformer-based methods, while it

is set to 50 for the LSTM architecture. In Table 2, we have also included additional experimental results from State Tennessee, which has a high level of SO_2 emissions based on the EPA reports [36], with prediction time window set as 24 hours. Table 2 has shown that Health-driven Opt based on transformer architecture achieves the lowest loss in health impacts compared to other methods.

Implementation Details In our experiments, predictions are made across different time window steps, denoted as T. We set T to values of 24 and 72 hours to explore the impact of varying prediction time windows. Temporal sequences are handled by slicing inputs and targets according to these specified sliding window steps T, and an 80/20 train-validation split is employed. For the CISO region dataset, we utilize the Stochastic Gradient Descent (SGD) optimizer with a learning rate of 0.001 and a batch size of 32. In addition to our Transformer-based models, we also implement LSTM-based variants of both the Fuel-mix-driven Opt and Health-driven Opt methods as **baselines** for comparison. These LSTM baselines use the same optimization objectives as their corresponding Transformer-based counterparts, i.e., Health-driven Opt and Fuel-mix-driven Opt, respectively. All experiments are conducted on a single NVIDIA K80 GPU. Training the Transformer-based models for 100 epochs takes usually less than five hours.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction claim to link electricity usage to public health using a domain-specific AI model for social well-being. We build an end-to-end pipeline, `HealthPredictor`, to address this, and validate its effectiveness through experiments across multiple U.S. regions, see Section 5.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: See section 7.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.

- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [NA]

   Justification: NA

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We include the main experiment setting and dataset source in the main body (Section 4, Section 5), and include all the information in the Appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Our datasets and code will be released upon publication of our paper.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The main body includes our model architecture, baseline model, loss function, and key hyperparameters, which are adequate for appreciating the results 5. More details can be found in the Appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Due to time and compute resource constraint (Section A.2), we do not conduct repeated experiments.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
   - If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: See Section A.2

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conform with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Our research mainly focus on positive societal impacts, currently we do not foresee any negative impacts of the work.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets we use are publicly available and mentioned with corresponding sources in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.